

Periodically-forced finite networks of heterogeneous globally-coupled oscillators: a low-dimensional approach

Carlo R. Laing^{a,*}, Ioannis G. Kevrekidis^b

^a *Institute of Information and Mathematical Sciences, Massey University, Private Bag 102-904 North Shore Mail Centre, Auckland, New Zealand*

^b *Department of Chemical Engineering and Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ 08544, USA.*

Abstract

We study a network of 500 globally-coupled modified van der Pol oscillators. The value of a parameter associated with each oscillator is drawn from a normal distribution, giving a heterogeneous network. For strong enough coupling the oscillators all have the same period, and we consider periodic forcing of the network when it is in this state. By exploiting the correlations that quickly develop between the state of an oscillator and the value of its parameter we obtain an approximate low-dimensional description of the system in terms of the first few coefficients in a polynomial chaos expansion. Standard bifurcation analysis can then be performed on the low-dimensional system which results from this computational coarse-graining, and the results obtained from this predict very well the behaviour of the high-dimensional system for any set of realisations of the random parameter. Situations in which the method begins to fail are also discussed.

Key words: Equation-free, coupled oscillators, bifurcation, polynomial chaos
PACS: 05.45.Xt, 05.45.a, 07.05.Tp, 82.40.Bj

1 Introduction

Synchronisation is a common phenomenon in biology and elsewhere [1,16,29,33]. It is often studied by investigating the conditions under which oscillators in a

* Corresponding author. Ph: +64-9-414-0800 x 41038. Fax: +64-9-441-8136.
Email addresses: c.r.laing@massey.ac.nz (Carlo R. Laing),
yannis@princeton.edu (Ioannis G. Kevrekidis).

particular network will synchronise [3,25]. Periodic forcing of systems is also ubiquitous [16,39], and so it is natural to study the entrainment of a network of coupled oscillators by a periodic forcing function. Many authors have studied small networks of two or three non-identical oscillators [3], and larger networks of oscillators that may have some symmetry [18,25] or a particular form of coupling [34]. The continuum limit in which there exists an infinite number of oscillators has also been studied in detail and many results are known for this case [2,4,17,24,38]. However, it is known that finite networks can show behaviour that does not occur in the continuum limit [4,12]. In many situations, finite networks are the most realistic way to model a physical system [6,12,35]. Results for large, finite networks will thus help bridge the gap between small network dynamics (for which bifurcation analysis is straightforward) and those for an infinite number of oscillators (where statistical physics provides the appropriate tools).

In this paper we consider a large but finite heterogeneous network of globally-coupled oscillators, which are collectively periodically forced. However, we do not analyse the system exactly; instead we analyse a low-dimensional description of it. This is the “equation-free” approach developed by Kevrekdis et al. [23]. The results here extend those of Moon et al. [30,32], obviating a reduction to phase oscillators: we consider two-variable (rather than single variable) oscillators, capable of undergoing Hopf bifurcations; we consider periodic forcing of the network, and we perform bifurcation analysis on the low-dimensional description of the system to understand how the behaviour of the full system changes as parameters are varied. We will not take the continuum limit as the number of oscillators tends to infinity, but instead analyse the realistic case of a finite number of oscillators.

The system we study is

$$\frac{dx_i}{dt} = y_i - x_i \left[x_i^2/3 - (\phi + \beta\mu_i) \right] + x_i^2/2 - \frac{\epsilon}{N} \sum_{j=1}^N (x_i - x_j) \quad (1)$$

$$\frac{dy_i}{dt} = -x_i + A \sin(\omega t) \quad (2)$$

for $i = 1, \dots, N$, where N is the number of oscillators in the network. For most of this paper we set $N = 500$. The oscillators are van der Pol oscillators [19] with an extra term ($x_i^2/2$) which breaks the internal symmetry $[(x, y) \rightarrow (-x, -y)]$ of each individual oscillator. These oscillators were chosen as being “typical” in the sense of not having any particular properties, e.g. symmetry. For $\beta = 0$, an uncoupled oscillator ($\epsilon = 0$) undergoes a supercritical Hopf bifurcation as ϕ increases through zero, with angular frequency 1. The angular frequency for an isolated oscillator as a function of ϕ is shown in Fig. 1.

The μ_i are taken from a normal distribution with mean 0 and standard deviation 1. (As discussed below, the methodology can be used with other distributions.) If $\beta \neq 0$ the network is heterogeneous, and each oscillator, if uncoupled, would have a different angular frequency determined by the value of $\phi + \beta\mu_i$. When $A = 0$, for β small enough, ϕ of moderate size and ϵ large enough, the oscillators synchronise in the sense of having the same period. Note that oscillators i and j cannot synchronise in the sense of $x_i(t) = x_j(t)$ for all t unless $\mu_i = \mu_j$. In this synchronised state the attractor of the system is a periodic orbit, which could be parametrised by a periodic variable, say $\theta(t)$. The variables $x_1, \dots, x_N, y_1, \dots, y_N$ could each then be written as functions of θ . This description would no longer be valid if one or more of the oscillators “unlocked” from the group.

We want to study the system in this synchronised state, but do not want to keep track of all the $2N$ variables $x_1, \dots, x_N, y_1, \dots, y_N$. Instead, we describe the state of the system by a small number of variables. We cannot easily derive an equation that governs the dynamics of these variables, but by repeatedly mapping between the two levels of description of the system we can numerically evaluate the results of integrating these unavailable equations; we can also find their collectively periodic states and their dependence on parameters, without ever obtaining the reduced equations in closed form. This low-dimensional description results in computational savings by, for example, giving a much smaller Jacobian matrix.

If the system is in this synchronised state and we increase A from zero, it will become periodically driven and it may be possible for the oscillators to lock with the driving frequency [16]. The latter part of this paper will consider this phenomenon in detail, but we first discuss the particular low-dimensional description of the forced system (1)-(2) used here, and how it can be used in projective integration to speed up direct simulation of this system.

2 A low-dimensional description

The main idea behind the low-dimensional description used here depends on correlations that rapidly develop between x_i, y_i and the value of μ_i in the parameter regime where synchronization eventually prevails. This is demonstrated in Fig. 2 where we plot the x_i and y_i as functions of μ_i for a particular realisation of the μ_i at three different times. We see that after just two periods of the forcing strong correlations develop between the state of an oscillator and its μ_i value. We will see that these correlations occur whether or not the network is synchronised with the forcing.

These correlations allow us to expand the x and y in certain classes of poly-

nomials of μ [30]. Xiu and Karniadakis [40] showed that if the μ_i are normally distributed, Hermite polynomials are the best in the sense of giving the fastest convergence of the expansion. For other μ distributions, the expansion converges fastest for other polynomial types; for example, if the μ_i are taken from a uniform distribution then Legendre polynomials are appropriate, and if a gamma distribution is used then Laguerre polynomials are best. Using other polynomials corresponds to the so-called Generalised Polynomial Chaos, or GPC [40]. We write

$$x(t, \mu) = \sum_{j=0}^q a_j(t) H_j(\mu) \quad (3)$$

$$y(t, \mu) = \sum_{j=0}^q b_j(t) H_j(\mu) \quad (4)$$

where H_j is the j th Hermite polynomial [$H_0(x) = 1, H_1(x) = 2x, H_2(x) = 4x^2 - 2, \dots$]. This expansion is known as a polynomial chaos expansion [30], and the a_i and b_i are the polynomial chaos coefficients. For a specific realisation of the μ_i , we have $x_i(t) = x(t, \mu_i)$ and similarly for the $y_i(t)$. Our low-dimensional description then involves the $2(q+1)$ coefficients $a_0, \dots, a_q, b_0, \dots, b_q$. This description is approximate, and the approximation becomes better as q is increased. Given x_i and y_i for a particular set of μ_i , the a_i are found by minimising the quantity

$$\sum_{i=1}^N \left[x_i - \sum_{j=0}^q a_j H_j(\mu_i) \right]^2 \quad (5)$$

and the b_i are found by minimising

$$\sum_{i=1}^N \left[y_i - \sum_{j=0}^q b_j H_j(\mu_i) \right]^2. \quad (6)$$

This is easily done in Matlab using the “backslash” operator to solve an overdetermined linear system. The operator from the x_i and y_i to the a_i and b_i is referred to as the “restriction” operator. Similarly we construct a “lifting” operator: given the $a_i(t)$ and $b_i(t)$ and a particular realisation of the μ_i we have

$$x_i(t) = \sum_{j=0}^q a_j(t) H_j(\mu_i) \quad (7)$$

$$y_i(t) = \sum_{j=0}^q b_j(t) H_j(\mu_i) \quad (8)$$

Armed with these two operators we can now proceed to numerically solve the unavailable equation for the polynomial chaos coefficients.

3 Coarse Projective Integration

Coarse projective integration entails accelerating the simulation of a high-dimensional system by projecting forward in time using only the variables in a low-dimensional description of the system. This is accomplished by occasionally performing short bursts of full simulation of the high-dimensional system in order to obtain the numerical information (estimation of the time-derivatives of the low dimensional description variables) required to perform accurate projections [23,30,32]. We can use the low-dimensional description in the previous section for coarse projective integration as follows. For convenience, let the high-dimensional description be the variable

$$X = [x_1, \dots, x_N, y_1, \dots, y_N] \in \mathbb{R}^{2N} \quad (9)$$

and the low-dimensional polynomial chaos coefficient description be the variable

$$Z = [a_0, \dots, a_q, b_0, \dots, b_q] \in \mathbb{R}^{2(q+1)}. \quad (10)$$

Given $X(0)$, integrate (1)-(2) forward for N_1 steps of size δt . Calculate Z at some or all of the times $t = 0, \delta t, 2\delta t, \dots, N_1\delta t$ using the restriction operator. Use these values of Z to extrapolate the values of Z to a time $N_2\delta t$ in the future, i.e. to time $(N_1 + N_2)\delta t$. Lift from the value of $Z((N_1 + N_2)\delta t)$ to $X((N_1 + N_2)\delta t)$ as detailed above. Restart the integration of (1)-(2) using $X((N_1 + N_2)\delta t)$ as the initial condition and integrate for a further N_1 time steps. Restrict to Z and repeat the procedure. If the cost of restricting, extrapolating and lifting is small compared to the cost of integrating the system (1)-(2) for N_2 time steps, this procedure may well be faster than integrating (1)-(2) directly. We expect this to be the case when the full system is characterized by a separation of time scales; the same principle underpins several analytical reduction techniques (e.g. centre manifolds, inertial manifolds) but in our case the reduction is obtained on the fly, from the short full simulation bursts.

We show results in Figs. 3 and 4 for $\delta t = 0.005, N_1 = 3$. In this case the projective step involves fitting a cubic polynomial to 4 data points (the last 3 of which were obtained through direct integration of the full system) and then evaluating this polynomial at a time $N_2\delta t$ in the future. Explicitly, if the four values of the j th component of Z obtained from integration of the full system

are $(Z_{-3}^j, Z_{-2}^j, Z_{-1}^j, Z_0^j)$, we fit the cubic $f^j(t) = a_3^j t^3 + a_2^j t^2 + a_1^j t + a_0^j$ through the points $(-3\delta t, Z_{-3}^j)$, $(-2\delta t, Z_{-2}^j)$, $(-\delta t, Z_{-1}^j)$ and $(0, Z_0^j)$. The extrapolated value of Z^j is then $f^j(N_2\delta t)$. The top panel of Fig. 3 shows the speedup as a function of N_2 . The speedup is defined as the time taken to directly integrate (1)-(2) over $0 < t < 100$ with time-step δt divided by the time taken to integrate over $0 < t < 100$ using coarse projective integration, as described. A speedup greater than 1 (N_2 greater than approximately 10) means that projective integration is more efficient than direct integration (provided that the accuracy in the values of the reduced variables is satisfactory).

Of course, as N_2 is increased the integration will start losing accuracy. The bottom panel of Fig. 3 shows the results of an integration when $N_2 = 1$ and the two curves shown (one for projective integration and one for direct integration) are indistinguishable. The top panel of Fig. 4 shows the case when $N_2 = 71$. In this case, coarse projective integration involves taking 3 steps of length δt , giving the clusters of 4 points shown in the bottom panel of Fig. 4, and then projecting the a_i and b_i forward a time $71\delta t$, lifting these values to initialise the x_i and y_i and continuing. The integration is clearly less accurate than that shown in Fig. 3, but the general behaviour is still qualitatively reproduced. In the spirit of Taylor series approximations, the extrapolation can only be accurate up to some fixed interval into the future, so as δt is increased, N_2 must be decreased, and the speedup will decrease (of course, the accuracy of the full integration will then also decrease). The results shown here will change if N_1 is changed or a different extrapolation scheme is used. A full analysis of projective integration for the system discussed here is beyond the scope of this paper (see discussions in [13,15,23,36,37]); it is clear, however, that for problems with a large separation of time scales and for appropriate parameter choices, it will be more efficient than straight integration. Step adaptation techniques from traditional numerical analysis based on *a posteriori* error estimates can be modified for the adaptive selection of projective steps.

Note that the simulations shown in Figs. 3 and 4 started at $t = 0$, and thus show transient behaviour, and different initial conditions were used for the two simulations. It is also important to note that we chose different realisations of the μ_i at each lifting step; the results are therefore representative of the expected behaviour *over different realizations* of the random variable. Should we only be interested in the acceleration of computations for a particular, single realization, the results would be even more accurate.

4 The 1:1 orbit

Consider the case of 1:1 locking, i.e. solutions for which each oscillator undergoes one oscillation during each forcing period. The usual way to study this would be to “strobe” the system once each forcing cycle. Defining X_p to be the state of the system at $t = 2\pi p/\omega$, where p is an integer, i.e.

$$X_p = X(2\pi p) \in \mathbb{R}^{2N} \quad (11)$$

where X is defined in (9), we could construct a map $g : \mathbb{R}^{2N} \rightarrow \mathbb{R}^{2N}$ as

$$X_{p+1} = g(X_p) \quad (12)$$

A 1:1 locked orbit is then a fixed point of g and its stability is determined by the eigenvalues of the Jacobian of g , evaluated at the fixed point. However, finding such a fixed point by, for example, Newton’s method, is computationally very expensive due to the high dimensionality of the system. Also, the results we obtain will only be correct for the particular realisation of the μ_i , a point to which we return below. Instead we use the low-dimensional description of the system in terms of the variable $Z \in \mathbb{R}^{2(q+1)}$. Defining $Z_p = Z(2\pi p)$, where Z is defined in (10), we can construct a map $h : \mathbb{R}^{2(q+1)} \rightarrow \mathbb{R}^{2(q+1)}$ as

$$Z_{p+1} = h(Z_p) \quad (13)$$

(From now on we choose $q = 1$, so $h : \mathbb{R}^4 \rightarrow \mathbb{R}^4$.)

Note that a fixed point of h is generally not a fixed point of g ; however, we will see that fixed points of h do describe the overall behaviour of the system, and the stability follows from the eigenvalues of the Jacobian of h , evaluated at its fixed points. In Fig. 5 we show the difference in x values (and in y values) after a time of one period, for a fixed point of h , i.e. a 1:1 locked orbit in the variables a_0, \dots, b_1 . We can see that none of the oscillators returned precisely to its initial condition. The two distributions $x_i(0)$ and $x_i(2\pi/\omega)$ give the same values of a_0 and a_1 , even though they clearly do not completely coincide. Similarly for $y_i(0)$ and $y_i(2\pi/\omega)$. If the order of our approximation (i.e. q) was increased, the discrepancy shown in Fig. 5 would decrease and the fixed point of h would better approximate the fixed point of g .

To evaluate $h(Z)$ in practice, we lift from Z_p to X_p using (7)-(8), integrate (1)-(2) for one period, then restrict from X_{p+1} to Z_{p+1} using (5)-(6). Although we could use projective integration as described in Sec. 3 to integrate (1)-(2), for simplicity we did not use projective integration while obtaining any of the

following results. We find the Jacobian of h using finite difference approximations to the derivatives, but we could also use matrix-free methods such as Newton GMRES [22] to find fixed points of h .

5 Continuation

We can continue fixed points of h as parameters in (1)-(2) are varied using standard pseudo-arclength continuation software [9]. In Fig. 6 we show the 1:1 locked orbit as ω is varied for a single oscillator (or equivalently, the network with $\beta = 0$, since in this case all oscillators behave identically). The left and right boundaries of the closed curves are saddle-node bifurcations where stable and unstable 1:1 locked orbits annihilate one another [8].

We want to analyse the case when $\beta \neq 0$, i.e. when the network is heterogeneous. We could do this for a single realisation of the μ_i as above, but to be more general we choose a number of different realisations of the μ_i and average over them. We do this averaging within our definition of the map h . Suppose r is the number of realisations we average over. For each $j = 1, \dots, r$ we calculate $Y_{p+1}^j = h(Y_p)$ using the j th realisation of the μ_i . (Note that Y_p is fixed.) We then define the averaged map \hat{h} as

$$\hat{h}(Y_p) = \frac{1}{r} \sum_{j=1}^r Y_{p+1}^j. \quad (14)$$

The results of implementing this averaging and following the 1:1 orbit are shown in Fig. 6 (dashed line). We can see that the effect of the heterogeneity is to move the range of ω values for which there is locking to lower frequencies. Even though the behaviour of the system was determined by following fixed points of \hat{h} , the results agree extremely well with those found from direct numerical integration of the full system (1)-(2) for any realisation of the μ_i from the correct distribution.

5.1 Varying A

We can follow the saddle-node bifurcations of Fig. 6, which mark the edges of the locking region, as both A and ω are varied. The results are shown in Fig. 7, where the resonance ‘‘tongues’’ for a single oscillator (solid line) and a network of 500 oscillators with $\beta = 0.5$ (dashed line) are shown. This figure shows more clearly that the effect of the heterogeneity is to move the 1 : 1 locking region to lower frequencies. Fig. 6 is a horizontal slice along the top boundary of Fig. 7.

5.2 Varying β

5.2.1 Breakdown of the reduced description

It is clear that increasing β increases the heterogeneity of the network. To understand the effects of this, in Fig. 8 we plot the boundaries of the 1:1 tongue as ω and β are both varied. We see the tongue boundaries move to lower frequencies, as expected from previous results. (Note that Fig. 6 shows slices through Fig. 8 at $\beta = 0$ and $\beta = 0.5$.) When the system is unforced, β and ϵ act in opposition: if the heterogeneity (i.e. β) is increased, the coupling strength (ϵ) must be increased in order to keep the network synchronised. However, we consider ϵ to be fixed. Thus for β large enough the forced system will no longer act as a “super-oscillator” in which all of the oscillators are synchronised with each other. Once this occurs the concept of locking between all oscillators and the forcing signal is no longer valid and the algorithm for following “coarse” (or macroscopic) saddle-node bifurcations terminates due to a lack of convergence within user-specified tolerances. We demonstrate this phenomenon in detail in Fig. 9.

The top two panels of Fig. 9 show the behaviour for a typical realisation of the μ_i just outside the 1:1 tongue, for a high value of β . The oscillators are ordered by their μ values. The behaviour of the 10 oscillators with highest μ_i is shown in panel A. In this case, oscillators 1-492 are synchronised with each other, but oscillators 493-500 are not synchronised with the rest of the group. However, oscillators 1-492 have also lost their locking to the forcing signal, and this is demonstrated on panel B, where we plot x_1 as a function of time. This slow (apparently quasiperiodic) modulation is typical for an oscillator just outside a resonance tongue. (Plotting x_i for any $1 \leq i \leq 492$ would give a similar picture.)

Panels C and D show the behaviour just inside the tongue (note: a different realisation of the μ_i from that in panels A and B has been used). Here, oscillators 1-496 are synchronised with each other, but oscillators 497-500 are not synchronised with the first 496. However, now oscillators 1-496 appear to be still, for all practical purposes, entrained by the forcing. This is shown in panel D, where x_1 is plotted as a function of time. The (apparently) periodic oscillation shown here has the same frequency as the forcing, and a plot of x_i for any $1 \leq i \leq 496$ would be very similar. We say “apparently” periodic motion because once one oscillator has desynchronised from the main group, none of the oscillators will undergo truly periodic motion. Instead, the motion is expected to be quasiperiodic with at least two frequencies present, or maybe even weakly chaotic.

For the results shown in Fig. 9 (with $\beta = 1.2$) approximately 1% of the oscil-

lators (those with the highest values of μ) are not synchronised with the main cluster, either inside or outside of the tongue. However, the remaining $\sim 99\%$ are synchronised with each other and using the “macroscopic” approach taken here we can detect whether this large cluster is synchronised with the forcing signal or not.

As β is increased, the fraction of oscillators no longer locked to the main cluster increases and the description of the system from the macroscopic point of view as a forced super-oscillator, using polynomial chaos coefficients, becomes increasingly flawed. This is the reason for deciding to terminate the curves in Fig. 8. Note that the two curves in Fig. 8 terminate at different values of β , but for both curves, the saddle-node bifurcation following algorithm fails to converge within tolerances when approximately 1% of the oscillators become desynchronised from the main group.

Note that increasing the number of Hermite polynomials, q , used in the macroscopic description (thus increasing the accuracy of the low-dimensional description) will not allow these curves to be followed to greater values of β . It is the lack of synchrony within the forced network that underlies the termination of the curves. Of course, increasing ϵ would allow the curves in Fig. 8 to be meaningfully continued to higher values of β .

Note that if we were to follow a vertical path through the middle of the tongue shown in Fig. 8 for a particular realisation of the μ_i , there would be many “fine-scale” bifurcations as one or more oscillators desynchronised from the main group. However, these are not visible in our macroscopic description of the system; we would need to change our macroscopic description in order to detect them [31].

Moon et al. [30,32] also considered the loss of synchrony in a heterogeneous network of Kuramoto oscillators. They were studying projective integration and showed that if one or two oscillators broke from the main cluster, projective integration could continue, as long as the low-dimensional description was augmented by the phase angle(s) of the oscillators that had lost synchrony. We take a different approach here, regarding the unsynchronised oscillators as providing a perturbation to the dynamics of the synchronised group.

5.2.2 Phase walkthrough

For a single periodically driven oscillator, “phase walkthrough” can occur just outside a 1:1 resonance tongue [14]. In this phenomenon the driven oscillator appears to be nearly synchronised with the driving oscillator, but every so often it undergoes either one extra or one fewer oscillation than the drive before returning to near synchrony. A similar phenomenon can occur in a pair of coupled oscillators. [5]. This is because the system lies in the vicinity of a

saddle-node bifurcation of periodic orbits. This walkthrough occurs approximately periodically, and the period scales as $|\omega - \omega^*|^{-1/2}$, where ω^* is the value of ω at the relevant tongue boundary [11]. As can be seen, this slow oscillation can be made arbitrarily slow by adjusting ω .

A similar phenomenon occurs in our system, but with a slight difference. For small β all of the oscillators are synchronised with one another, effectively acting as one oscillator, and we can observe phase walkthrough near the tongue boundaries with the scaling just mentioned above. However, this phenomenon is a result of the system spending a long time in phase space near the remains of the stable and unstable fixed points of \hat{h} , and is thus sensitive to noise or other perturbations.

Once at least one “runaway” oscillator has become desynchronised from the rest (as a result of increasing β) the system can be thought of as a noisily perturbed oscillator, the “oscillator” being the vast majority of oscillators that are synchronised with each other, and the “noise” resulting from the influence of the desynchronised oscillator(s) on the rest. Thus we expect that we can no longer make the slow oscillation arbitrarily slow just by adjusting ω . Indeed, for large fixed β , near the boundaries shown in Fig. 8 there is a range of ω values for which the slow oscillation (walkthrough) period is not well-defined, since perturbations from the desynchronised oscillator(s) affect the neutrally stable behaviour at the underlying bifurcation, resulting in apparently stochastic “slipping” relative to the forcing signal. It may be possible to describe these rare occurrences in terms of Langevin dynamics on a low-dimensional free energy surface [20,27].

5.3 Varying ϕ

Another parameter of interest to vary is ϕ . Recall that varying ϕ in a single unforced oscillator causes a Hopf bifurcation, leading to oscillations. The result of varying ϕ is shown in Fig. 10, for both a single oscillator and for a network with $\beta = 0.3$. We see that the tongue terminates at a positive value of ϕ , and that heterogeneity moves the tongue boundary to lower values of ϕ .

To understand the cusps for low values of ϕ we plot in Fig. 11 a cross-section through Fig. 10 at $\phi = 0.8$, for a single oscillator. The four saddle-node bifurcations are clear. (For the network, a similar plot is found, not shown.) The cusps involve both saddle-node bifurcations being annihilated at a codimension-two point. In the vicinity of these cusps, previous results on the periodically forced van der Pol oscillator [19] show that there should be a curve of Hopf bifurcations of the fixed point of the map starting near each cusp, which will be associated with the generation of quasiperiodic motion. We can follow these

curves using standard algorithms [10], and the results for the left cusp are shown in Fig. 12, both for a single oscillator and for the inhomogeneous network of 500 oscillators. (We also followed the Hopf bifurcation curve associated with the right cusp, not shown.) The Hopf bifurcations correspond to a complex conjugate pair of eigenvalues crossing out of the unit circle in the complex plane as ω is decreased. Writing these eigenvalues at bifurcation as $e^{\pm i\theta}$, we have $\theta = 0$ at the rightmost point of the Hopf bifurcation curve (i.e. eigenvalues of $+1, +1$) and θ monotonically increases as ω is decreased until $\theta = \pi$ (i.e. eigenvalues of $-1, -1$) at the leftmost point on the Hopf bifurcation curve.

Note that while following the curve of Hopf bifurcations for the network, we could not use a larger value of β than $\beta \approx 0.3$ (i.e. we could not use a more heterogeneous network) because for larger values of β the oscillators with the largest values of μ would become desynchronised from the rest as the bifurcation was approached. The problem discussed in Sec. 5.2 regarding the effectiveness of the macroscopic approach would then reoccur.

The Hopf bifurcation for a single forced oscillator is supercritical, with a stable 2-torus being created as the Hopf bifurcation curve is crossed in the direction of decreasing ω [19]. The criticality of the Hopf bifurcation for the network seems to be the same as that for a single oscillator, and even though the curve in Fig. 12 was found by averaging over 20 realisations of the μ_i , it is still a very good predictor of the parameter values at which quasiperiodic dynamics occur for any particular realisation of the μ_i (not shown).

Previous results [19] lead us to expect (for each cusp) a curve on which there are orbits homoclinic to the fixed point of the map (i.e. homoclinic to a periodic orbit in the full system), emanating from the point where the curve of Hopf bifurcations and saddle-node bifurcations meet, for both a single oscillator and the network. We do not consider these curves further.

6 Discussion

In this paper we studied a finite network of heterogeneous, globally-coupled oscillators, all subject to the same periodic forcing. We coarse-grained the dynamics, obtaining a low-dimensional description of the system in terms of a few polynomial chaos coefficients. We defined a return map by sampling the low-dimensional system once every forcing period; by finding and following fixed points of this map we performed standard bifurcation analysis on the 1:1 locked state at a much reduced computational cost compared with analysis of the full system. By averaging over realisations of the distribution of the heterogeneity we have been able to obtain results valid for any particular realisation, rather than just a single one.

The main result of this paper is the demonstration that the dynamics of large, heterogeneous networks of forced coupled oscillators can be simulated and analysed using a low-dimensional description of the state of the network, provided that all or most of the oscillators are synchronised with one another. The reduction to a low-dimensional description is performed “on the fly” and results in considerable computational simplification. For this specific network and heterogeneity we found that the effects of heterogeneity are to move the 1:1 locked orbit to lower frequencies, and to lower values of ϕ .

We have concentrated on only the 1:1 resonance; the techniques used can be easily applied to any other resonances. One issue we have not discussed is varying N , the number of oscillators in the network. We found that for small values of N the number of realisations, r , of the μ_i that are averaged over in the definition of \hat{h} (eqn. (14)) had to be increased in order for continuation algorithms to converge to a given tolerance. This makes sense since, as N is increased, the difference between a simulation with one particular realisation of the μ_i and that for a simulation with a different realisation, will decrease, and thus fewer realisations will need to be averaged over.

Compared to the Kuramoto model results of Moon et al. [30,32], or their animal flocking models [31] where each oscillator consists of a single variable (a phase angle) the coupled units here are representative of general ODE-based oscillators, capable of undergoing Hopf bifurcations. It is this extra feature that gives rise to results such as those in Figs. 10 and 12.

It would be interesting to apply these ideas to networks of oscillators where the onset of oscillation is through different types of bifurcations, e.g. homoclinic bifurcation, saddle-node-on-a-circle (or SNIPER), or a saddle-node bifurcation of periodic orbits, or to other finite heterogeneous networks that have previously been studied [6,35]. Another possibility is to study the periodic forcing of networks of coupled bursting neurons. The response of isolated bursting neuron models to periodic forcing has recently been studied [7,26], as has the behaviour of coupled bursters [21].

Regarding the problems caused by one or more oscillators desynchronising from the main group (see Sec. 5.2), one way to deal with this might be to expand the states of oscillators in functions other than “globally” defined polynomials of μ . For example, a wavelet basis of localised functions may be more suitable [28], particularly if the network breaks into clusters, with the oscillators within each cluster being synchronised.

Acknowledgements: We thank Sung Joon Moon for useful conversations about the work presented here, and the referees for their useful comments. The work of IGK was partially supported by DARPA and DOE.

References

- [1] Acebrón, J. A., Bonilla, L. L., Vicente, C. J. P., Ritort, F. and Spigler, R. The Kuramoto model: A simple paradigm for synchronization phenomena. *Rev. Mod. Phys.* **77**, pp. 137-185, 2005.
- [2] Ariaratnam, J. T. and Strogatz, S. H. Phase diagram for the Winfree model of coupled nonlinear oscillators. *Phys. Rev. Lett.* **86** pp. 4278-4281, 2001.
- [3] Baesens, C., Guckenheimer, J., Kim S. and MacKay, R. S. Three coupled oscillators: mode-locking, global bifurcations and toroidal chaos. *Physica D.* **49** pp. 387-475, 1991.
- [4] Balmforth, N. J. and Sassi, R. A shocking display of synchrony. *Physica D.* **143**, pp. 21-55, 2000.
- [5] Borgers, C. and Kopell, N. Effects of noisy drive on rhythms in networks of excitatory and inhibitory neurons, *Neural Computation* **17** pp. 557-608, 2005.
- [6] Butera, R. J., Rinzal, J. and Smith, J. C. Models of respiratory rhythm generation in the pre-Bötzinger complex. II. Populations of coupled pacemaker neurons. *J. Neurophysiol.* **82** pp. 398-415, 1999.
- [7] Coombes, S., Owen, M. R. and Smith, G. D. Mode-locking in a periodically forced integrate-and-fire-or-burst neuron model. *Phys. Rev. E*, **64**, 041914, 2001.
- [8] Devaney, R. L. An introduction to chaotic dynamics systems, second edition. Addison-Wesley, 1989.
- [9] Doedel, E. J. AUTO: Software for Continuation and Bifurcation Problems in Ordinary Differential Equations, 1997. <http://indy.cs.concordia.ca/auto/>
- [10] Doedel, E., Keller, H. B. and Kernevez, J. P. Numerical analysis and control of bifurcation problems (I) Bifurcation in finite dimensions. *Int. J. Bifn. Chaos*, **1** (3) pp. 493-520, 1991.
- [11] Doiron, B., Laing, C., Longtin A. and Maler, L. Ghostbursting: a novel neuronal burst mechanism. *J. Computational Neuroscience*, 12(1), pp. 5-25, 2002.
- [12] Doiron, B., Rinzal, J. and Reyes, A. Stochastic synchronization in finite size spiking networks. *Phys. Rev. E.* **74**, 030903, 2006.
- [13] E, W. and Engquist, B. The heterogeneous multiscale methods. *Comm. Math. Sci.* **1**, pp. 87-132, 2003.
- [14] Ermentrout, G. B. and Rinzal, J. Beyond a pacemaker's entrainment limit: phase walk-through. *Am. J. Physiol. Regulatory Integrative Comp. Physiol.* **246**, pp. 102-106, 1984.
- [15] Gear, C.W. and Kevrekidis, I.G. Projective methods for stiff differential equations: problems with gaps in their eigenvalue spectrum. *SIAM J. Sci. Comput.* **24**, pp. 1091-1106, 2003.

- [16] Glass, L. Synchronization and rhythmic processes in physiology. *Nature*. **410**, pp. 277-284, 2001.
- [17] Golomb, D., Hansel, D., Shraiman, B. and Sompolinsky, H. Clustering in globally coupled phase oscillators. *Phys. Rev. A*. **45**, pp. 3516-3530, 1992.
- [18] Golubitsky, M., Stewart, I., Buono, P. L. and Collins, J. J. Symmetry in locomotor central pattern generators and animal gaits. *Nature*. **401**, pp. 693-695, 1999.
- [19] Guckenheimer, J. and Holmes, P. *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer-Verlag, New York. 1990.
- [20] Haataja M., Srolovitz D. J. and Kevrekidis I. G. Apparent Hysteresis in a Driven System with Self-Organized Drag. *Phys. Rev. Lett.* **92**, 160603, 2004.
- [21] Izhikevich, E. M. Synchronization of elliptic bursters. *SIAM J. Appl. Math.* **60**, pp. 503-535, 2000.
- [22] Kelly, C. T. *Iterative methods for linear and nonlinear equations*. SIAM, 1995.
- [23] Kevrekidis, I. G., Gear, C. W., Hyman, J. M., Kevrekidis, P. G., Runborg, O. and Theodoropoulos, C. Equation-free, coarse-grained multiscale computation: enabling microscopic simulators to perform system-level analysis. *Comm. Math. Sci.* **1**, pp. 715-762, 2003.
- [24] Kuramoto, Y. *Chemical Oscillations, Waves, and Turbulence*, Springer, Berlin, 1984.
- [25] Laing, C. R. Rotating waves in rings of coupled oscillators. *Dynamics and stability of systems*. **13**, pp. 305-318, 1998.
- [26] Laing, C. R. and Coombes, S. Mode-locking in a periodically forced “ghostbursting” neuron model. *Int. J. Bif. Chaos*, **15**, pp. 1433-1444, 2005.
- [27] Laing, C. R., Frewen, T. and Kevrekidis, I. G. Coarse-grained dynamics of an activity bump in a neural field model. Submitted, 2006.
- [28] Le Maître, O. P., Knio, O. M., Najm, H. N. and Ghanem, R. G. Uncertainty propagation using Wiener-Haar expansions. *J. Comp. Phys.* **197** pp. 2857, 2004.
- [29] Matthews, P. C., Mirollo, R. E. and Strogatz, S. H. Dynamics of a large system of coupled nonlinear oscillators. *Physica D*. **52**, pp. 293-331, 1991.
- [30] Moon, S. J., Ghanem, R. and Kevrekidis, I. G. Coarse-graining the dynamics of coupled oscillators. *Phys. Rev. Lett.* **96**, 144101, 2006.
- [31] Moon S. J., Nabet B., Leonard N. E., Levin S. and Kevrekidis I. G. Heterogeneous animal group models and their group-level alignment dynamics; an equation-free approach. *J. Theor. Biol.*, **246**, pp. 100-112, 2007.
- [32] Moon, S. J. and Kevrekidis, I. G. An equation-free approach to coupled oscillator dynamics: the Kuramoto model example. *Int. J. Bifn. Chaos*, **16** (7) pp. 2043-2052, 2006.

- [33] Pikovsky, A., Rosenblum, M. and Kurths, J. Synchronization. Cambridge University Press, Cambridge, 2001.
- [34] Ren, L. and Ermentrout, B. Phase locking in chains of multiple-coupled oscillators. *Physica D* **143**, pp. 56-73, 2000.
- [35] Rubin, J. and Terman, D. Synchronized activity and loss of synchrony among heterogeneous conditional oscillators. *SIAM J. Appl. Dyn. Syst.* **1**, pp. 146-174, 2002.
- [36] Samaey G., Kevrekidis I.G. and Roose D. Patch dynamics with buffers for homogenization problems. *J. Computational Physics* **213**, pp. 264-287, 2006.
- [37] Samaey G., Roose D. and Kevrekidis I.G. The gap-tooth scheme for homogenization problems. *Multiscale Modeling and Simulation*. **4** pp. 278-306, 2005.
- [38] Strogatz, S.H. From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators. *Physica D*, **143** pp. 1-20, 2000.
- [39] Winfree, A. T. The geometry of biological time. Springer-Verlag, 1990.
- [40] Xiu DB and Karniadakis GE. The Wiener-Askey polynomial chaos for stochastic differential equations. *SIAM J. Scientific Computing* **24**, pp. 619-644, 2002.




Fig. 1. Angular frequency of an isolated oscillator ($\beta = 0 = \epsilon$) as a function of ϕ .




Fig. 2. Left column: x_i as a function of μ_i (dots). Right column: y_i as a function of μ_i (dots). Also included are the polynomial chaos expansions with $q = 1$ (lines). From top to bottom: $t = 0, 2\pi/\omega, 4\pi/\omega$. Parameters are $A = 0.5, \omega = 0.85, \beta = 0.1, \epsilon = 1, N = 500$.

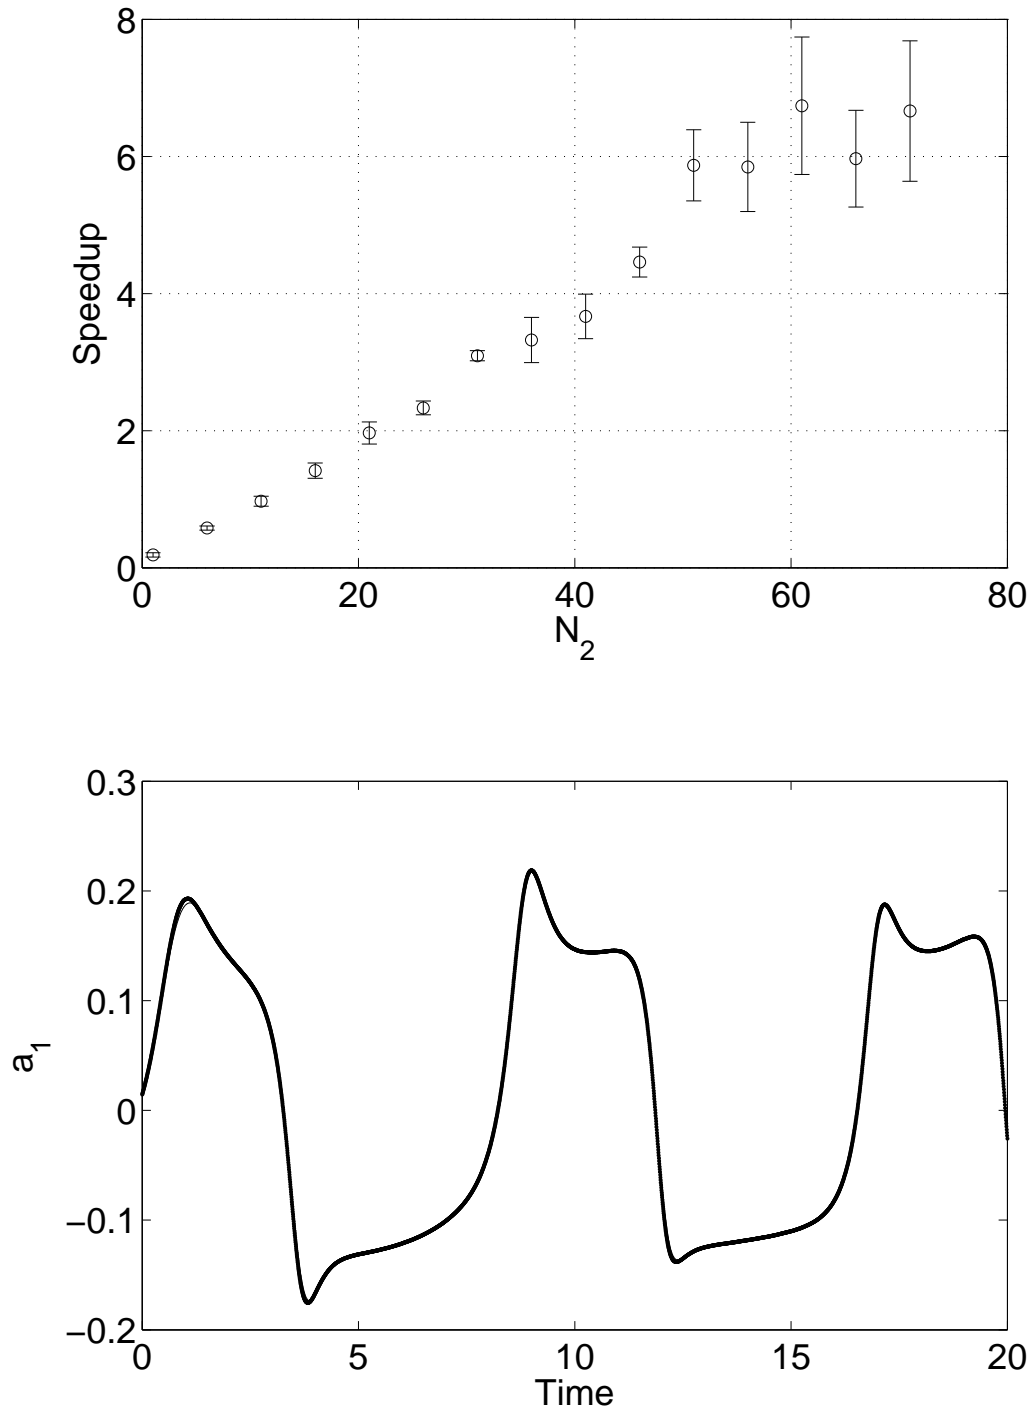


Fig. 3. Top: speedup (as defined in the text) as a function of N_2 , with $N_1 = 3$. Bottom: a_1 for projective integration with $N_2 = 1$ (dots) and a_1 from full integration (solid line — indistinguishable from the dots). Other parameters are $A = 0.5, \omega = 0.85, \phi = 1, \beta = 0.5, \epsilon = 1, q = 2$. See text for details.

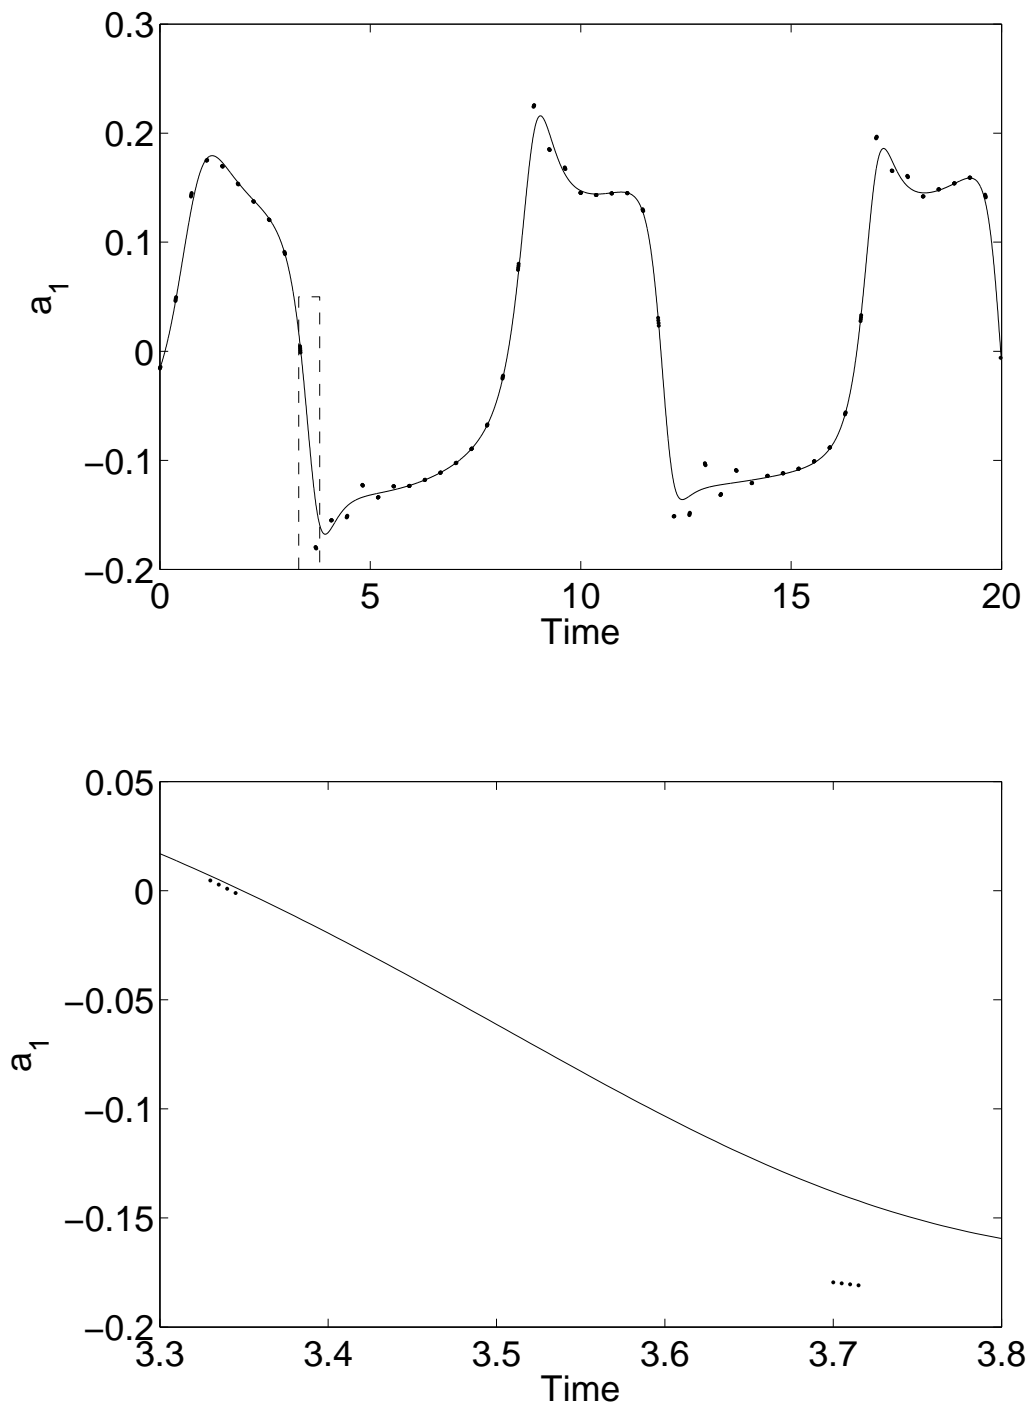


Fig. 4. Top: a_1 for projective integration with $N_2 = 71$ (dots) and a_1 from full integration (solid line). Bottom: blowup of the small rectangle in the top panel. Other parameters are $A = 0.5, \omega = 0.85, \phi = 1, \beta = 0.5, \epsilon = 1, q = 2$. See text for details.




Fig. 5. Difference in x values after one period (circles), and difference in y values (crosses), for a fixed point of h (i.e. for an orbit that is periodic in the coefficients a_0, \dots, b_1). Parameters are $A = 0.5, \omega = 0.85, \beta = 0.5, \epsilon = 1$.




Fig. 6. The 1:1 orbit. Solid line: x at multiples of $2\pi/\omega$ for a single oscillator. Dashed line: a_0 at multiples of $2\pi/\omega$ for a network with $N = 500$ and $\beta = 0.5$, averaged over $r = 20$ realisations. Other parameters are $A = 0.5, \epsilon = 1, \phi = 1$.




Fig. 7. Boundaries of the 1:1 Arnold tongue. Solid line: one oscillator, dashed: averaging over $r = 20$. $N = 500$ and $\beta = 0.5$. Other parameters are $\epsilon = 1, \phi = 1$.




Fig. 8. Boundaries of the 1:1 Arnold tongue, averaging over between 20 and 60 realisations. The curves terminate as β is increased because the oscillators become too heterogeneous to synchronise among themselves. $N = 500$ and $A = 0.5$. Other parameters are $\epsilon = 1, \phi = 1$.




Fig. 9. A and B: $\omega = 0.935, \beta = 1.2$ (just to the right of the right boundary shown in Fig. 8). C and D: $\omega = 0.925, \beta = 1.2$ (just to the left of the right boundary). Panels A and C show the evolution of x_{491} to x_{500} (colour coded, white high), while panels B and D show x_1 as a function of time. $N = 500$ and $A = 0.5$. Other parameters are $\epsilon = 1, \phi = 1$.




Fig. 10. Boundaries of the 1:1 orbit. Solid line: one oscillator. Dashed line: a network with $N = 500$ and $A = 0.5, \beta = 0.3, r = 20, \epsilon = 1$.




Fig. 11. x at multiples of $2\pi/\omega$ for the 1:1 orbit of a single oscillator, as ω is varied. Solid line: stable, dashed: unstable. Parameters are $A = 0.5, \phi = 0.8$. This figure is a horizontal slice at $\phi = 0.8$ through Fig. 10.




Fig. 12. Boundaries of the 1:1 Arnold tongue (solid and dashed lines, for a single oscillator and a network of $N = 500$ oscillators, respectively) and Hopf bifurcation curves (dots and circles joined by lines, for a single oscillator and a network of 500 oscillators, respectively). Parameters are $A = 0.5, \beta = 0.3, r = 20, \epsilon = 1$. Compare with Fig. 10.